

## *SAM (REX/Ops) Operations Meeting 2008.06.17*

Adam (typist), SteveW, Mike, Eric, SteveS, Art, Dave, Krzysztof, Anil, Robert, Angela, Randolph

Meeting time: 40 min

Posting these minutes resolved DZSAM-105

For reports submitted to the meeting...

*The report text is in this font*

*Comments made during the meeting are in this font inline with the report*

### **Releases:**



**sam\_db\_srv\_pkg\_v8\_4\_5** ( 13/Jun/08 | [Release Notes](#) )



**SAMDEV-13** **FIXED**

*When a station adds a new cached file, it incorrectly sets the uncaching work grp id instead of the owner work grp id*



**SAMDEV-7** **FIXED**

*Fix availability status so that people can better select a particular status*

*Above is in UPD. Can't install yet - hardware problem on bluearc servicing Kits server (permissions are wrong). SteveW will forcefully and personally open a helpdesk ticket.*

*SAMDEV-13 fixes a station hang for station with multiple groups. (we think this was introduced in 8.4.1).*

*SAMDEV-7 is Randolph's fix.*

*Angela will be adding to integration.*


*Need to fix the DB with this release.*

*Upcoming releases...*



**sam\_db\_srv\_pkg\_v8\_4\_6** ( Release Notes )
 Progress: 
 0 of 1 issues have been resolved


[SAMDEV-12](#) **UNRESOLVED** [Change D0 autodeestination logic for -test versions](#)



**sam\_db\_srv\_pkg\_v8\_5\_0** ( Release Notes )
 Progress: 
 0 of 2 issues have been resolved



[SAMDEV-6](#) **UNRESOLVED** [samadmin remove cached files command does not copy the cached files record to the history table](#)


[SAMDEV-9](#) **UNRESOLVED** [Incorrect exception specification for DataFileImpl.removeReplica\(\)](#)


**sam\_cpp\_api\_v8\_4\_2** ( Release Notes )
 Progress: 
 1 of 1 issues have been resolved








[SAMDEV-4](#) **FIXED** [Incorrect user check in c++ client](#)


**sam\_v8\_4\_2** ( Release Notes )
 Progress: 
 1 of 1 issues have been resolved


[SAMDEV-11](#) **FIXED** [SAM client not properly threadsafe](#)

*Fixed means its in CVS, but not released into Kits yet.*

*Issues not assigned to a release yet*

T	Key	Status	Summary
	<a href="#">SAMDEV-5</a>	 Open	Improve project robustness against DB failures
	<a href="#">SAMDEV-1</a>	 Open	Convert station to use omniORB
	<a href="#">SAMDEV-10</a>	 Open	Problem with --route in station v6_0_5_24_srm

**DØ:**

Things worth nothing:

-- [d0ols.fnal.gov](http://d0ols.fnal.gov) --- dead for for quite >3 days --- Is this "critical node" going to be backed up soon ? Or will be like that for a long time.

*If just one or two are down, then that's considered ok. Adam will ask online group to*

*find out what is going on.*

-- June 15 (night) : DDE (*dataset definition editor*) webpage down due to hung tomcat process hanging on to the required network port and preventing the new process from starting correctly. Robert fixed it by manually finding and killing it. Shifters don't/didn't know how to do this.

-- June 15/16: D21CLTO Mover was dead (RED status).

Unresolved issues (initiated by users):

DZSAM-100 : MC Jobs overwhelming SBGRID --- It seems to me that this ticket could be closed as 2 GB tgz file is a feature of D0 software and there have been cleaning scripts, as Adam said.







DZSAM-62 : "40 MC jobs held over the weekend" in OSG Grid  
--- Is this still a problem ? Working on it ?










*Robert - Nameservice died last night possibly due to disk problems on d0ora2. Anil says there is a failure, SAN and admins are looking into it. Backups are failing - i/o errors on SAN. DB is on a different, more expensive (and hopefully more reliable) SAN.*

*d0srv071 still has disk problems. (do ls on certain files in disk cache on two different partitions get i/o error, but other things are ok). Will be rebooted this am (hopefully no fsck). Rebooting did not work. Will need to fsck.*

*d0srv072 filled up - no log file cleanup script (there is one now).*

*osgouhep station - context service created a 335GB trace file on d0srv066. Ask Andrew about this. Perhaps the culprit for Corba comm failures. (red on SAAG). Just added DZSAM-104 for this issue.*

Issues: <b>DZSAM resolved issues from past week</b> (Displaying 6 of 6)			
	<a href="#">DZSAM-102</a>	fcg daemons on d0srv072 have crashed because / is 100% full	
	<a href="#">DZSAM-99</a>	D0 control room says that they are getting luminosity DBserver errors	
	<a href="#">DZSAM-97</a>	Please install new sam_osg_handler.sh on OSG forwarding nodes	
	<a href="#">DZSAM-96</a>	Samgrid job not submitted all the OSG jobs to cab server	
	<a href="#">DZSAM-94</a>	Post minutes from 2008.06.10 SAM Ops Meeting	
	<a href="#">DZSAM-71</a>	Add stale globus job manager reaper to crontab	

Issues: <b>DZSAM unresolved problems</b> (Displaying 5 of 5)			
	<a href="#">DZSAM-104</a>	OSG-OUHEP Context server big log and Corba comm failures	
	<a href="#">DZSAM-100</a>	MC Jobs overwhelming SBGRID	
	<a href="#">DZSAM-62</a>	40 MC jobs held over the weekend	
	<a href="#">DZSAM-54</a>	Jobs don't work on Lancaster Forwarding Node	
	<a href="#">DZSAM-53</a>	samgrid reporting incorrect status of jobs	

## CDF:

End of CDF SAM Shift June 10-17 2008 summary

fcdfdata320 (4) offline (since months ago)

fcdfdata154 (4 read pools) offline (since before last week)

*Hardware problem - motherboard was replaced, but still does not come up. We think it is still under warrantee.*

Not pingable, email sent to helpdesk. (cdf\_dcache\_help was not working)

fcdfdata157 (3 write pools) was offline, now OK *Angela restored.*

r-fcdfdata146-1 r-fcdfdata146-1,2,3,4Domain OFFLINE since Sunday.

*dCache admins contacted, but no response yet. Angela was not able to bring it back up herself.*

Email sent to cdf\_dcach\_help about all 8 down pools.

Enstore bad mover:

9940B22 offline ("replace system disk"--since before last week)

Tues June 10 15:00

Angela restarted the DBServers using new DBaccounts--

It turned out many job were hung, and needed to be  
restarted later.

Wed June 11, 14:00

fcdfdata135-1 went offline ; Angela fixed it around 15:30

Wed June 11, 14:00

LT02\_29 bad mover--fixed 17:00

IAB386 had a bad label--replaced Thur June 12 11:50

by Stanley Hicks.

While IAB386 was unavailable, there were lots of jobs hung

trying to access it (up to 70) which showed up at

DCache "difference between requested and queued"

Friday:

Corrupted files reported on dCache pool r-fcdfdata097-1. Files have been removed.

Reference JIRA ticket CDFDH-129. (Angela)

Monday evening: Total number of files to restore very large:1966

/3 difference. Sent email to dcache-admin.

Tuesday 17<sup>th</sup>:

DCache Door 21 has 177 logins. (red)

No.of files to restore/ diff requested & queued 2717

*Angela says no action needed.*

CSL to Enstore status:

Tue Jun 17 07:10:45 2008 1213704645 Warning:






























b0csl21 -- Clsam & Clmeta may be dead

























*Problem with last modified time of the log file --- the log is being updated, but the last modified file time is not updated by nfs. So this confused monitoring. Investigating. Angela is interacting with the helpdesk. Looking into what gets affected when it reboots. Also looking at fixing monitoring scripts to key off of something else instead of the last modified time. See CDFDH-120.*

CDF-LT03 Some red/orange reads ongoing & pending.

*Red/orange even though things are working. Requests are backed up.*

Issues: **CDFDH resolved issues from past week** (Displaying 15 of 15)

	<a href="#">CDFDH-129</a>	Corrupt file on Linux fcdldata097.fnal.gov 2.6.9-34.ELsmp #1 SMP Tue Mar 14 10:51:18 CST 2006 i686 i686 i386 GNU/Linux	
	<a href="#">CDFDH-125</a>	Re: unexplained rise in difference between requested and queued	
	<a href="#">CDFDH-124</a>	LT03_29 bad mover, IAB386 bad tape	
	<a href="#">CDFDH-119</a>	Diskpool door not responding	
	<a href="#">CDFDH-118</a>	Bad Mover LTO3_04 offline	
	<a href="#">CDFDH-117</a>	cell offline r-fcdldata097-1	
	<a href="#">CDFDH-114</a>	Fermigrid jobs failing because of null pointer to dataset.	
	<a href="#">CDFDH-111</a>	SAM failures	
	<a href="#">CDFDH-108</a>	[Fwd: failed CAF job from 5/27 afternoon - STKEN problem?]	
	<a href="#">CDFDH-107</a>	r-fcdldata135 misbehaving again	
	<a href="#">CDFDH-105</a>	/var nearly full on fcdldata108	
	<a href="#">CDFDH-103</a>	Bad Mover/Bad Tape	
	<a href="#">CDFDH-88</a>	Re: DCache Cell offline - fcdldata154	
	<a href="#">CDFDH-58</a>	DCache machine fcdldata157 offline	
	<a href="#">CDFDH-12</a>	Re: error accessing SAM database	

Issues: CDFDH unresolved problems (Displaying 12 of 12)			
	<a href="#">CDFDH-128</a>	cdf-caf and cdf-fncdfrsv0 stations can get stuck if restarted	
	<a href="#">CDFDH-127</a>	data access problem	
	<a href="#">CDFDH-121</a>	Debugging SAM @ CNAF	
	<a href="#">CDFDH-120</a>	Clsam and clmeta on b0csl21 will not stay alive	
	<a href="#">CDFDH-113</a>	user strolog/cdf/namcaf@fnal.gov is not authorized	
	<a href="#">CDFDH-112</a>	re : fcdfrdata135 is offline	
	<a href="#">CDFDH-109</a>	SAM problems @ CNAF	
	<a href="#">CDFDH-106</a>	data handling, stuck processes	
	<a href="#">CDFDH-96</a>	cdf-japan station is down	
	<a href="#">CDFDH-55</a>	Re: copying files to Bologna	
	<a href="#">CDFDH-30</a>	Inactive Dcache cells	
	<a href="#">CDFDH-23</a>	Stuck dCache requests	

## MINOS:

Minos 'Week In The Woods' annual meeting last week in Ely, MN.

Enstore -

DCache -

Jun 24 FNDCA public DCache upgrade to 1.8 - is shifted to Jun 24

We will do some testing before Jun 24.

Jun ? FNDCA public DCache downtime for SDE all day

SAM -

SAM-IT/3582 - provided login and testing instructions to developers.

This failed again with --constrain-delivery removed.

Waiting for further advice.

Can I access [samread@fcdfrsam1.fnal.gov](mailto:samread@fcdfrsam1.fnal.gov) to look at CDF configs ?

Oracle -

Oracle quarterly patches - scheduled Jun 24 during DCache upgrade.



## OPEN ISSUES

IT 1146: cannot test existence of sam tape location

IT 2071: quadratic project startup time growth

PLAN -

## DB:

D0 :

- NFS problem on SAN and home areas on d0ora2 ~8:15PM on Monday 06/18

Got the errors like "SCSI transport failed"

"write error on /dev/dsk/c3t210000D023600221d0s0"

- DB is fine. However, backup failed last night due to the error. Seems like home area cleared itself.

SAN storage admin and sysadmin for d0ora2 will be investigating the issue.

Owner	Object	Type	Extents Used	Extents Allocated	Extents Left
SAMPRD	CACHED_FILE_PROJ_USAGE_HISTS	TABLE	96	100	4

Action: Doubled the maxextents

Information for: d0ofprd1

Owner	Object	Type	Extents Used	Extents Allocated	Extents Left
SAMPRD	PF1_UQ1_UK	INDEX	3097	3100	3

Action : Increased the maxextents to 6,100

CDF :

Owner	Object	Type	Extents Used	Extents Allocated	Extents Left
SAMPRD	PROC_AF_FK_I	INDEX	996	1000	4

Action: Doubled the maxextents

Alert(s):

MINOS:

- Scheduling Oracle's Quarterly patch on minosprd June 24th at 8:30AM

## **dCache/enstore:**

enstore/dcache report

-----

Scheduled Downtimes:

- o stken - **CANCELLED** (*this was scheduled for 6/19, not enough personnel*)
  - public enstore s/w upgrade to SDE

- o fndca - Tue, Jun 24, 8 am - 4 pm
  - public dcache s/w upgrade to release 1.8

- o d0en - Tue, Jul 1 (??), 7:30 am - 11:30 am
  - ADIC AML/2 maintenance

Operations:

- o d0en

- o cdfen

- o cdfdca

- o stken
  - On Wednesday, all the LTO4 drives and some LTO3 and 9940B drives were offline for about an hour and a half due to the accidental renaming of a couple of open files. The files were found and restored.

- o fndca